Nonvolatile RAMs arranged as 256 × 4 and 64 × 4 bits facilitate the design of microprocessor-based systems requiring permanent storage of such parameters as calibration data for instruments.

4-bit-wide nonvolatile RAMs open up new applications

With the appearance of nonvolatile RAMs in 4-bit-wide configurations, designers can now enjoy the benefit of these unusual memories while cutting the chip count. These parts are unique in that they contain two memories in one device, a RAM and an EEPROM (see "How Nonvolatile RAMs Work").

The RAM functions as a fully static memory. The EEPROM, on the other hand, can be accessed only through the RAM. The storage operation transfers the entire contents of the RAM to the EEPROM with a single TTL pulse of 100 ns or greater. The recall operation transfers the contents back to the RAM with a 450-ns TTL pulse. At the end of either operation, both memories will contain the same data. The block diagram of Xicor's nonvolatile RAM family is shown in Fig. 1.

The first nonvolatile RAM was organized as 1024 × 1 bit. Although this configuration is extremely useful, eight chips are needed to construct a byte-wide memory. Some applications, however, require less than 1 kbyte of memory, and for them the designer had to buy the additional capacity even though he did not need it.

Since then, two more parts have come out, a 64 × 4 and a 256 × 4 device. The main advantage of the smaller part is in the storage of data that change infrequently—for instance, toggle-switch settings in terminals or calibration data in instruments. The larger-capacity memory is more applicable to those uses where data generated during on-line operation must be retained when power is removed.

To illustrate how these new ×4 memories can be incorporated into a design, it is helpful to look first at the interface signals of the memory and how they interact with a host computer.

Figure 2 shows the pin configuration of the three chips. The 254 × 4 memory, the X2212, comes in an 18-pin DIP. The X2210, organized as 64 × 4 bits, is pin-compatible with the X2212, with two unused pins for the higher-order addresses. The earlier part, the X2201, has been given a new command structure, and the new version is also being introduced with the new parts.

The new command structure makes the nonvolatile RAM considerably easier to use. Three changes have been made: Only one pulse is required for a storage operation; a $\overline{CS}$ signal is no longer required for either storage or recall; and storage is blocked with a recall-low—that is, the recall line is low prior to the occurrence of the storage pulse.

The use of a single storage pulse allows the microprocessor to issue a storage command and then perform other tasks without any further involvement in the storage operation. Removing $\overline{CS}$ as a requirement for storage and recall simplifies the logic to generate the $\overline{CS}$ signal. Finally, storage with a recall-low signal is a more positive way to ensure against an inadvertent storage, since it is difficult to guarantee a high from a TTL gate while the power supply is rising.

To demonstrate the easier-to-use interface of the

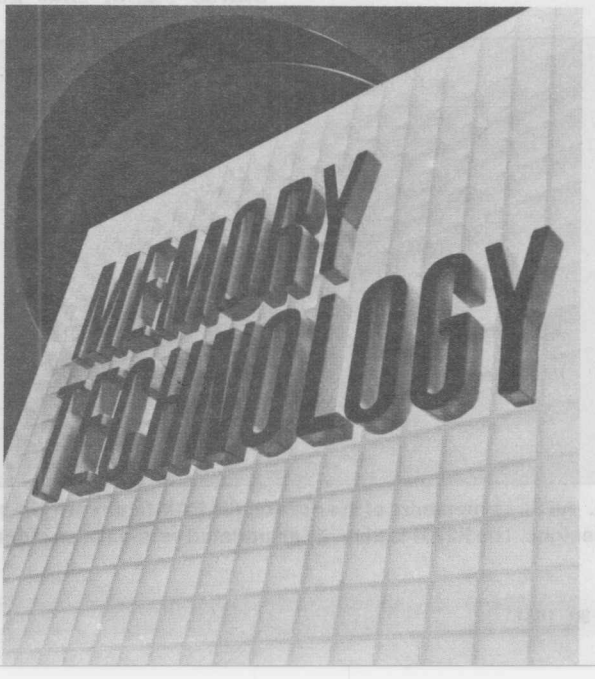

George Landers,
Applications Manager
Xicor, Inc.
851 Buckeye Ct.
Milpitas, CA 95035

new command structure, consider the X2210 used with an Apple computer. The Apple II was selected because its address space is fully decoded for I/O functions and it has convenient card slots using these decoder outputs. In fact, three separate decoder lines run to each card slot. One is common for all card slots and is 2048 bytes located at hexadecimal address $C800 through $CFFF; the other two are exclusive for an individual card slot. For slot 6 these are a 16-byte block starting at hex address C0E0 and a 256-byte block starting at $C600.

Adding nonvolatile RAMs to a computer

When the Apple II is turned on, it looks for a disk controller card by examining the first 8 bytes of the 256-byte block for each card slot. When the correct 8 bytes are found, the computer jumps control to the program located in the 256-byte block.

By locating the CS signal for the RAM at $C600 and loading the correct first 8 bytes, the Apple II can be fooled at power-on into turning control over to a program in the nonvolatile RAM. The recall line is tied to the reset of the Apple, and the storage line (labeled Store) to the 16-byte exclusive select at $C0E0. Figure 3 shows that no components other than the two X2210s are needed for the interface.

Figure 4 is a machine-language program that draws a box on the screen, with the byte at $C61B determining the location of the upper left-hand corner and byte $C613 the location of the lower right-hand corner. When the Apple is turned on and the program is placed in slot 6, the box will automatically be displayed on the screen. The corners can be positioned by writing $C613 or $C61B and then reading $C0E0. The latter operation initiates storage, so that the size of the box will be different when the computer is turned on again.

The preceding descriptions assume that the power supply of the system is clean in its power-up and power-down sequences. In most systems, that may not be the case and should not be assumed. The nonvolatile RAM inhibits any operations whenever the power supply is below 3 V, but some form of storage protection should be provided above 3 V.

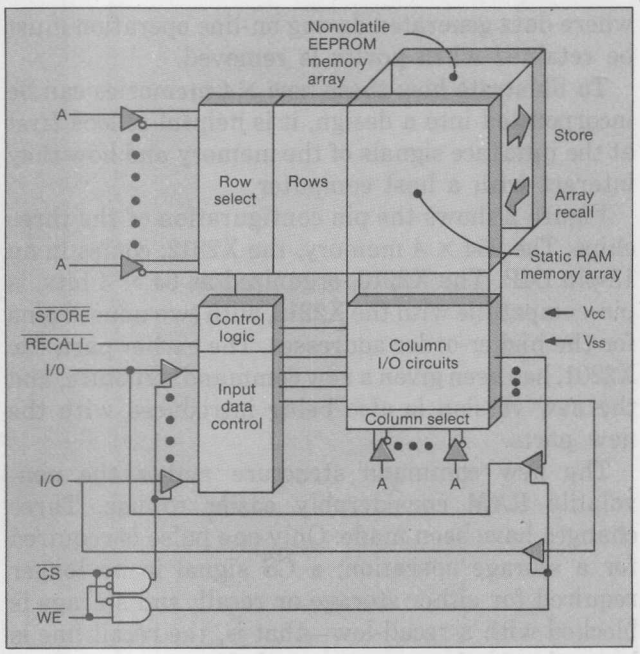

1. The secret behind a nonvolatile RAM is an on-chip EEPROM that can hold the entire contents of the RAM.

Storing data permanently

There are two methods of storage protection, and at least one or preferably both should be designed into the system. The two basic actions are to ensure that the voltage on the storage pin follows the power-supply voltage or to hold the recall pin low while the power-supply voltage is moving up or down. The latter can be accomplished with a power-supply status signal that goes low whenever the power

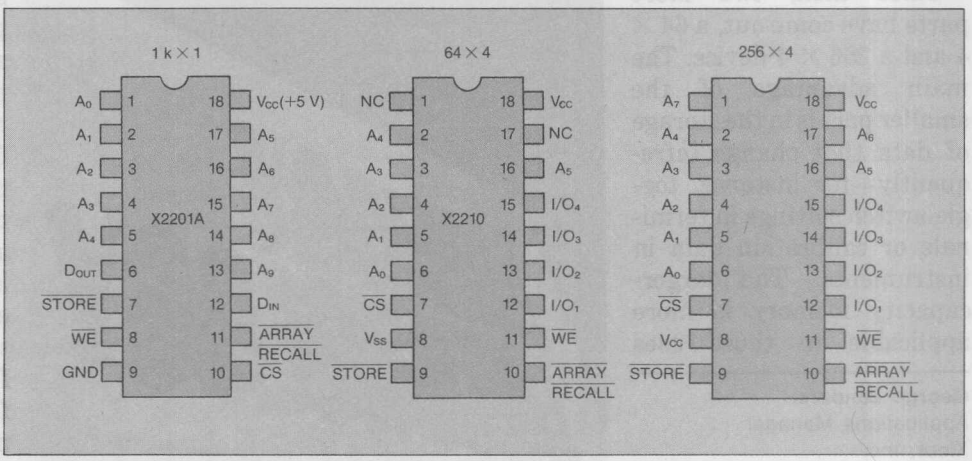

2. All three members of the nonvolatile RAM family are pin-compatible, easing many new designs. The X2201 is actually an updated version of an older part.

supply is out of tolerance. Tying this signal to the input of the open-collector NAND gate in Fig. 5 ensures that the voltage level on the storage pin trades the supply and that the recall pin is held low. An open-collector gate is used, since the output transistor cannot receive base current as long as any input is held low.

Some of the schemes employed to provide the power-supply status signal involve zener diodes and a sense circuit to monitor either the regulated or the unregulated dc. Alternatively, optoelectronic ac fault detectors can monitor the ac line. The circuit shown in Fig. 6 is a simple monitor of the regulated dc. Whenever the power supply drops below 4.5 V, transistor Q_1 loses base current and turns off, turning Q_2 on. The capacitor provides some delay before releasing the power-supply status.

Example shows how

Figure 7 shows an example of a design using storage protection techniques. The storage pin is held high whenever one of the inputs to the NAND gate is low. For a storage command to be received, several conditions must exist: the power supply must be above 4.5 V, the system reset must be high, the system recall must not be selected, the system write must be low, and the system storage line must be selected. Any other condition and the open-collector NAND gate output will not turn on.

The R/$\overline{W}$ line is included as a gate input, since during initiation, the microprocessor can generate a write command until it is fully under control. The recall line is pulled low whenever the power supply is below 4.5 V, the system reset is low, or the system recall is selected.

In this case, the memory card is designed to fit in slot 7 of the Apple II, with a disk controller card in slot 6. The decoder half selected by $C0F0-$C0FF is used to decode the storage and recall command. The system store line is selected when hex addresses $C0F0 through $C0F3 are written or read. Because of the ORing of the system write command, the nonvolatile RAM receives only a storage command for a write cycle. The system recall is selected for hex addresses $C0F4 through $C0F7 and can be either a read or a write cycle.

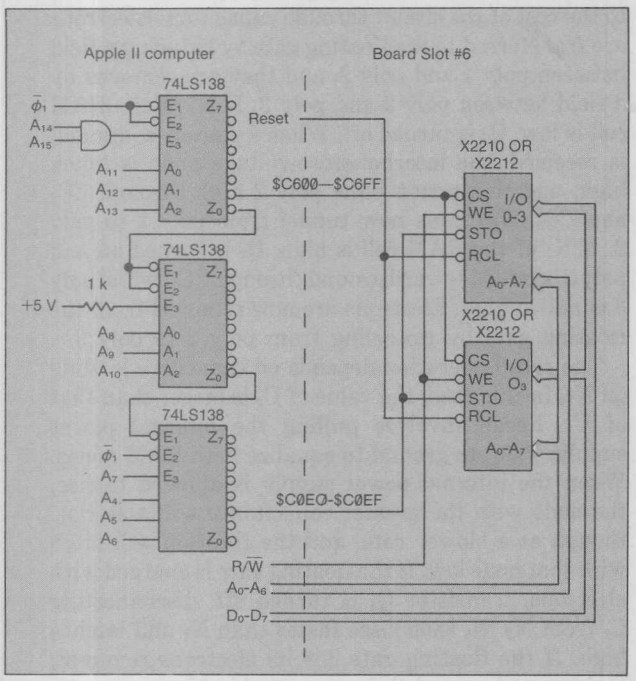

3. To operate the nonvolatile RAM on the I/O bus of an Apple II computer, the CPU reset line is connected to the memory's recall line and the address decoding signal for hex addresses C600 and C6FF selects the two memories.

Fooling the Apple

Again, if the first 8 bytes of the nonvolatile RAM located at hex address $C700 are correct, the Apple II will think that the card is a disk controller and give up control to the program it contains. A 1280-byte machine-language program could be written that does some tasks on power-up. This program can be changed at will by loading a new program from disk and issuing a storage command to place it in

```
*C600.C63F

C600- A2 20 A6 00 A2 03 86 3C
C608- 20 58 FC 20 40 F8 A9 FF
C610- 85 30 A9 27 85 2C 85 2D
C618- 85 00 A9 00 85 01 A4 01
C620- 20 19 F8 A4 01 A5 00 20
C628- 19 F8 A4 01 A5 01 20 28
C630- F8 A4 00 A5 01 20 28 F8
C638- A0 00 A1 00 20 28 F8 60
```

4. Shown here is a machine-language program that demonstrates how the nonvolatile RAM works with an Apple II computer. The program simulates the code for a disk controller card and then draws a box on the screen.

the EEPROM. Now, the computer performs a different task when it is turned on regardless of what disk is in the disk drive.

One novel use of nonvolatile RAMs has been to replace DIP switches as holders of configuration data or calibration constants. Here, opting for the memory reduces costs, improves reliability and density, and allows the designer to add system features. Another advantage is that the system board can be sealed against RFI, yet permit access to the switches (bits in the memory).

One memory equals 16 switches

As few as 16 switches (two packs of 8) can be replaced with a single X2210 at a lower overall—that is, assembly as well as component—cost. For a total of $5, the user now gets 256 switches. Direct savings also result from the elimination of hardware access. In the past, many systems provided access doors cut into the cabinet so that the user could change the settings on the switches. With the nonvolatile RAM, the "switches" can be accessed from a keyboard, and no doors, hinges, or clips are required.

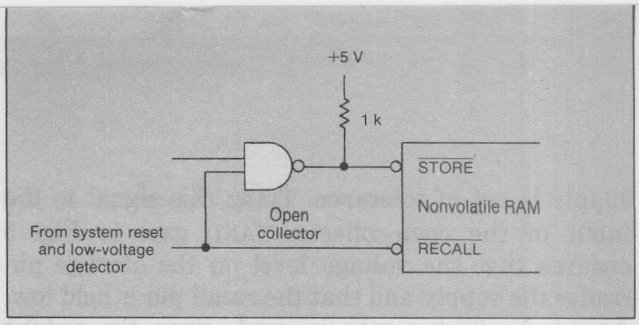

5. The open-collector NAND gate ensures that the storage pin is high whenever the power-supply status is low regardless of the power-supply voltage.

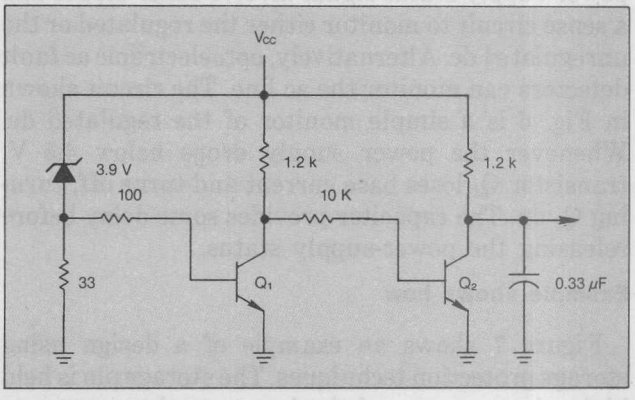

6. Forming a simple monitor of regulated dc, this circuit provides a low signal whenever the supply is below 4.5 V.

How nonvolatile RAMs work

The nonvolatile RAM consists of a RAM and an EEPROM on the same chip. The RAM operates as a fully static memory such as the popular 2102A or the 2114. The EEPROM stores or removes its data by the absence or presence of electrons on a floating gate. Except for the method of introducing and removing electrons—which does not require high-voltage or ultravoltage light sources—its operation is similar to that of UV EPROMs.

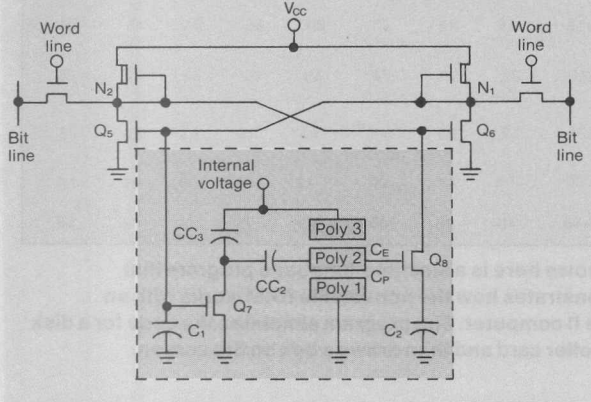

For its nonvolatile RAM, Xicor uses Fowler-Nordheim tunneling, a technique in which tunneling is enhanced through the use of textured surfaces that increase the local field strength. This approach obviates the use of very thin oxides to generate the necessary field strengths.

The figure shows the circuit diagram of each cell in the memory. The floating gate (poly 2) is connected to the rest of the circuit through capacitors. Electrons are transferred to the floating gate by inducing a field between poly 1 and poly 2, and they are removed by a field between poly 2 and poly 3. If N_1 of the RAM cell is low, Q_7 is turned off. When a storage command is received, the internal-store-voltage node is taken high, and this action pulls poly 2 high through CC_2 and CC_3. Electrons now tunnel from poly 1 to poly 2. If N_1 of the RAM cell is high, Q_7 is turned on and poly 2 is pulled toward ground through CC_2, while poly 3 is pulled high. Electrons are now removed from the floating gate by tunneling from poly 2 to poly 3.

The recall operation depends on capacitive loading on N_2. In this case, the value of C_2 is larger than that of C_1. Recall involves pulling the internal power supply, V_{CCA}, to ground to equalize both RAM nodes. When the internal power supply is allowed to rise, the node with the greater capacitance will also rise, though at a slower rate, and the flip-flop will latch with that node low. If the floating gate is charged with electrons, transistor Q_8 is turned off, disconnecting C_2 from N_2. N_2 then rises faster than N_1 and latches high. If the floating gate has its electrons removed, Q_8 is turned on and C_2 is connected to N_2, causing N_2 to rise slowly and latch low.

In addition, improvements in reliability usually yield cost savings from reduced rework, faster manufacturing throughput, and fewer returns. Nonvolatile RAMs benefit from the general reliability of semiconductors, as opposed to that of mechanical DIP switches. Furthermore, they are easier to test than DIP switches and often—since some DIP switches cannot be wave-soldered or placed in cleaning systems—easier to assemble and clean.

As mentioned, using nonvolatile RAMs, the system designer can add features to a system that would have been impossible with DIP switches. A system can now be reconfigured remotely by telephone link or from the keyboard with software prompting messages. If a system does not require many switches, the leftover part of the memory can be used to store nonvolatile operating data. Finally, the RAM can be used independently after the system has been configured at power-on.

When replacing DIP switches, data are seldom changed and the nonvolatile RAM functions as a read-mostly memory. Similarly, some systems designers are turning to nonvolatile RAMs to store changeable programs or firmware that can be downloaded over a remote line.

The other major use of this type of memory is to capture critical data in the event of a power failure. Examples of the sort of data required after power failure are system status, special accounting information, error conditions, and accumulated counts of events.

In these cases, the required procedure is simple. The best results occur when the status of the power supply is sensed ahead of the regulated dc signal—that is, either at the incoming ac or at the unregulated dc. When the sensed voltage is missing or low for a selected time, an interrupt is issued that sends a single TTL pulse to the storage pin of the nonvolatile RAM. The only requirement for the device is that the 5-V power supply remain within specification for at least 10 ms. During that time, the memory is off the bus and the microprocessor can take care of whatever orderly shutdown procedure is required by the system.□

How useful?	Circle
Immediate design application	562
Within the next year	563
Not applicable	564

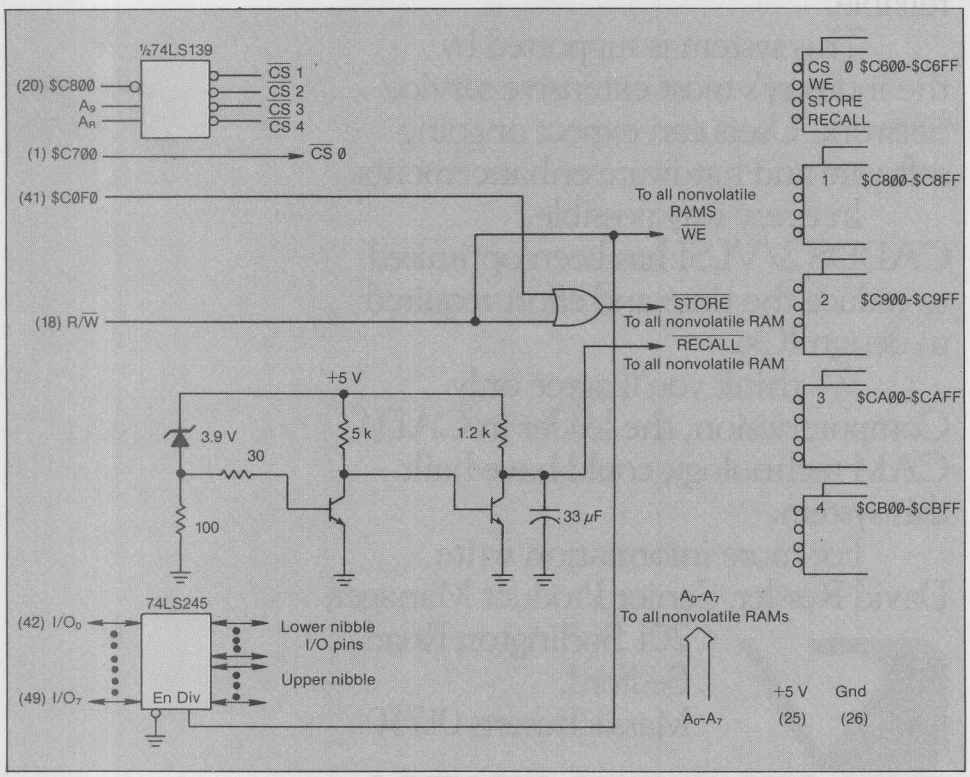

7. The circuit above connects nonvolatile RAMs to an Apple II bus and protects stored data on power-up and power-down.

Integrated circuits have to be complicated. Designing them does not.

CADDS™ 2/VLSI will speed the design and manufacture of even the densest ICs.

It's an extraordinarily productive system, combining very fast interaction with unsurpassed ease of use.

CADDS 2/VLSI has grown out of years' experience in the IC field. Like its predecessor, CADDS 2/VLSI is accurate, flexible, and thoroughly reliable.

This system is supported by the industry's most extensive service network. Users can expect ongoing software and hardware enhancements.

In every way possible, CADDS 2/VLSI has been optimized to reduce the time and effort required to design ICs.

We think you'll agree only Computervision, the leader in CAD/CAM technology, could have built this system.

For more information write David Ressler, Senior Product Manager,

201 Burlington Road,
Bedford,
Massachusetts 01730.

COMPUTERVISION

We see a world transformed.